arXiv:1609.00799v1 [cs.IR] 3 Sep 2016

Lexical-Morphological Modeling for Legal Text Analysis

Danilo S. Carvalho[1⋆], Minh-Tien Nguyen[1,2], Chien-Xuan Tran[1], and Minh-Le Nguyen[1]

[1] School of Information Science,
Japan Advanced Institute of Science and Technology (JAIST),
1-1 Asahidai, Nomi, Ishikawa, 923-1292, Japan.
[2] Hung Yen University of Education and Technology (UTEHY), Hung Yen, Vietnam.
{danilo, tiennm, chien-tran, nguyenml}@jaist.ac.jp

Abstract. In the context of the Competition on Legal Information Extraction/Entailment (COLIEE), we propose a method comprising the necessary steps for finding relevant documents to a legal question and deciding on textual entailment evidence to provide a correct answer. The proposed method is based on the combination of several lexical and morphological characteristics, to build a language model and a set of features for Machine Learning algorithms. We provide a detailed study on the proposed method performance and failure cases, indicating that it is competitive with state-of-the-art approaches on Legal Information Retrieval and Question Answering, while not needing extensive training data nor depending on expert produced knowledge. The proposed method achieved significant results in the competition, indicating a substantial level of adequacy for the tasks addressed.

1 Introduction

Answering legal questions has been a long-standing challenge in the Information Systems research landscape. This topic draws progressively more attention, as we experience an explosive growth in legal document availability on the World Wide Web and specialized systems. This growth is not accompanied by a matching increase in information analysis capabilities, which points to a severe under-utilization of available resources and to potential for information quality issues [1]. As a consequence, increasing pressure has been put into professionals of law, since having the relevant and correct information is a vital step in legal case solving and thus is closely tied to the matter of professional ethics and liability. This problem is often referred as the "information crisis" of law.

The ability to retrieve relevant and correct information given a legal query has improved over time, with the combination of expert Knowledge Engineering and Natural Language Processing (NLP) methods. However, the ability to answer questions in the legal domain is of special difficulty, due to the need of reasoning over different types of information, such as past decisions, laws and

⋆ Supported by CNPq – Brazil scholarship grant

facts. Furthermore, concepts in legal text are often used in a way that differs from common language use, and differences in laws and procedures from each country prevent the creation of comprehensive and coherent international law corpora. Common legal ontologies are among the efforts to facilitate automatic legal reasoning, but have not seen strong development in the past years [2]. In this context, *Textual Entailment Recognition* plays a very important role, as a set of hypothesis presented in a question will certainly have answers in the previously cited types of information (decisions, laws, facts). The Recognition of Textual Entailment (RTE) challenge series [3], although not specific to the legal domain, is a recognized benchmark for methods that can be adapted to legal texts.

To effectively answer legal questions, one fundamental set of information that must be available is the law, presented as the collection of codes, sections, articles and paragraphs that should be unequivocally referenced when a hypothesis is raised as part of a legal inquiry. Therefore, adequate representation of law corpora is the basis of a functional system for legal question answering. The representation problem is often associated with ontologies and other annotated knowledge bases, but these methods are costly and more difficult to automate when compared to fully text-based approaches, such as bag-of-words, n-gram and topic models.

In this work, we propose a fully text-based method for legal text analysis, in the context of the Competition on Legal Information Extraction/Entailment (COLIEE), covering both the tasks of Information Extraction and Question Answering. The goal is the retrieval of relevant law articles to a given yes/no legal question and the use of the retrieved articles to correctly answer the question in a completely automated way. Our contributions in this paper are as follows: (i) a ranking and selection method for legal information retrieval based on a mixed size n-gram model, including an original scoring function for ranking; (ii) an improved adaptation of a Textual Entailment classification method, based on Machine Learning ensembles (Adaboost), including a similarity feature built upon Distributional Semantics (Word2Vec). Lexical and morphological analysis were done on the English translated Japanese Civil Code, comprising tokenization, POS-tagging, lemmatization, word clustering and a set of lexical statistics. A study on success and fail cases is also provided, with common baseline practices and related works used as means of performance comparison. The results of COLIEE are presented as a means of substantiating the experimental evaluation and also discussing the proposed method's perceived shortcomings and improvements.

The remaining of this work is structured as follows: Section 2 presents the related works and relevant results; Section 3 details the Legal Question Answering problem and the COLIEE competition shared task; Section 4 explains our approach to the competition problem; Section 5 presents the experimental setting, results and discussion; Finally, Section 6 offers some concluding remarks.

[3] www.aclweb.org/aclwiki/index.php?title=Recognizing_Textual_Entailment

2 Related Works

Liu, Chen and Ho [3] presented the three-phase prediction (TPP) method for retrieval of relevant statutes in Taiwan's criminal law, given general language queries. The method was a hierarchical ranking approach to law corpora, featuring a combination of several Information Retrieval techniques, as well as Machine Learning and feature selection ones. Results were evaluated in terms of recall, achieving from 0.52 to 0.91, from the top 3 to 10 retrieved results, respectively.

Inkpen et al. showed one of the first successful models for RTE using SVMs [5]. Later, Castillo proposed a new system for solving RTE using SVMs [6], in which training data includes RTE-3, annotated data set from RTE-4, and the development set of RTE-5. 32 features were used and the training model achieved the best F-measure of 0.69 in two-way and 0.67 in three-way classification task.

Nguyen et al. [7] conducted a study of RTE on a Vietnamese version of RTE-3 [8] translated from Giampiccolo et. al. [9]. The author used SVMs trained with 15 features divided in two groups: distance and statistical features, in which the first group captures the distance and the second one represents the word overlapping between two sentences. A voting system combining three classifiers built on three feature groups (distance, statistical, and combined features) was used to judge entailment relation. The method obtained 0.684 of F-measure in two-way task.

In legal text, Tran et al. addressed legal text QA by using inference [10]. The author used requisite-effectuation structures of legal sentences and similarity measures to find out correct answers without training data and achieved 60.8% accuracy on 51 articles on Japanese National Pension Law.

Kim et al. proposed a hybrid method containing simple rules and unsupervised learning using deep linguistic features to address RTE in civil law [11]. The author also constructed a knowledge base for negation and antonym words which would be used for classifying simple questions. To deal with difficult questions, the author used morphological, syntactic and lexical analysis to identify premises and conclusions. The accuracy was 68.36% with easy questions and 60.02 with difficult ones.

This work uses all features in [7], as they apply to the same purpose. Additional features were also included: Word2Vec similarity and term frequency – inverse document frequency (TF-IDF). Our approach differs from [6] in using *Word2Vec*[17] similarity instead of *WordNet*.

3 Legal Question Answering

Legal Question Answering (LQA) consists in finding out and providing *"correct answers"* to a legal question given by users. An overview of LQA is shown in Fig. 1.

LQA can be divided in three tasks: 1) retrieving relevant articles, i.e., the ones containing the answer; 2) finding correct evidence in the relevant articles that allows answering the question; and 3) answering the question. While the first task is a specific case of Information Retrieval (IR), the second can be considered as a form of Recognition of Textual Entailment (RTE), in which given a question,

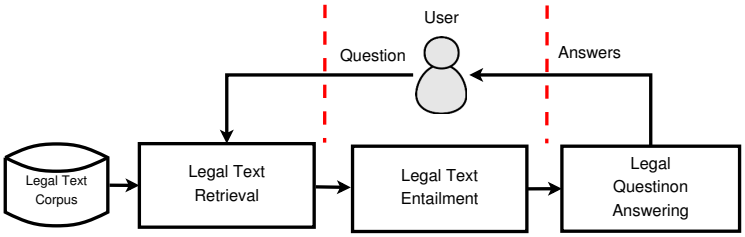

Fig. 1: The model of legal text question answering system

the LQA system has to decide whether and how a relevant article can answer the question. The third one is the final result of the two previous tasks, combined with answer formatting.

Legal text is considerably different to other types of text, e.g., news articles, due to their structural and semantic characteristics. Firstly, they have specific logical sentence structures e.g., requisite and effectuation [12]. Secondly, words and writing style are used in a strict form, because law documents require high correctness and should avoid ambiguity. Another aspect is that law documents are written in a high abstraction level [13]; therefore, they often require collection and linking of multiple concept references to enable understanding and answering of a question. The use of concept references leads to a situation in which there are few, or in some cases, even no overlapping words between a law question and its relevant articles.

In this work, LQA tasks are considered into the context of COLIEE, a competition on legal information extraction/entailment which was first held in 2014, in association with Workshop on Juris-informatics (JURISIN). COLIEE 2015 [4] is the second competition, consisting of three phases:

- Phase One: retrieving relevant articles from all Japanese Civil Code Articles given a set of YES/NO questions.
- Phase Two: evaluating the entailment relationship between the question and retrieved articles.
- Phase Three: combination of Phase One and Phase Two, the system will retrieve list of relevant articles given a query, and then decide the entailment relationship between retrieved articles and the provided question.

The Japanese Civil Code is composed by a collection of numbered articles, each one containing a set of declarations pertaining to a specific topic of the law, e.g., labor contracts, mortgages.

Information Retrieval Task: Relevance Analysis

The first phase consists on an explicit IR task, for which the goal is to retrieve the relevant articles that can be used to correctly answer a given yes/no question. The challenge in this task is to determine the relative relevance, i.e., Relevance Analysis (RA), of an article to the query presented in the question. Different

[4] webdocs.cs.ualberta.ca/~miyoung2/COLIEE2015/

articles dealing with the same topic often have similar wording and it is common for questions not to refer to topic keywords or refer to alternative ones. Furthermore, the restricted size of the Japanese Civil Code means that obtaining reliable linguistic information from articles is difficult and most questions will present new language structures that can range from useful to necessary for answering.

Simple Question Answering Task: Textual Entailment

The goal of Textual Entailment (TE) is to decide whether a legal query/question can be answered by a set of relevant articles retrieved with RA. This task can be accomplished by recognizing textual entailment (RTE), in which the query/question is treated as an hypothesis and relevant articles as evidence. Given a question Q and a set of relevant articles A, ($A = \{a_1, ..., a_n\}$), if Q is answered by a_i ($1 \leqslant i \leqslant n$), then a_i entails Q [9], [14]. A pair *(Q, a_i)* is assigned label YES if a entailment relationship exists, i.e., a_i can answer Q; otherwise, NO.

4 Proposed Approach

In order to be able to perform both Relevance Analysis and Textual Entailment recognition independently in phases one and two, and jointly in phase three, IR and classifier methods were developed separately. First, both the legal corpus and training data are analyzed and combined into representation models. The models are then used to rank articles or classify answers according to the task. The representation model used for Relevance Analysis is a mixed size n-gram collection and the one used for textual entailment are feature vectors for Machine Learning. Figure 2 shows the overall view of the proposed method.

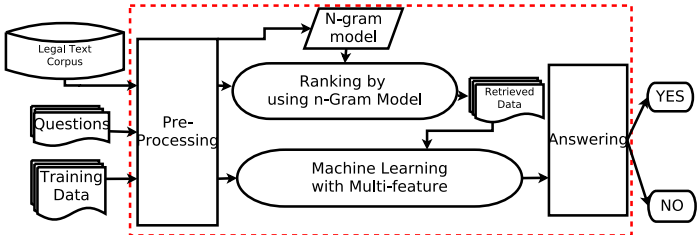

Fig. 2: Model overview

4.1 Relevance Analysis

A detailed analysis of the Civil Code and training data revealed that lexical and syntactic overlapping may vary to a high degree between questions and articles, and also between articles concerning the same topic. However, certain morphological features, such as lemmas, retain a higher level of consistency among topics. For this reason, the adopted representation model was a mixed size n-gram model, with $n : [1, k]$, i.e., terms made by sequences up to k words, in which the terms are lemmatized. For simplicity, the Relevance Analysis method

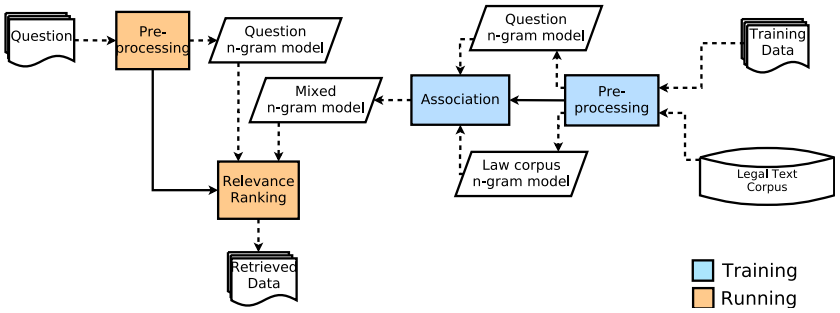

Fig. 3: The process of legal text retrieval

hereon described was named R_2NC (Ranking Related N-gram Collections). A summarized view of the process is shown on Figure 3.
The steps to build the model are detailed as follows:

1. Collect the entire content for each article, including section title;
2. Check references between articles and annotate accordingly;
3. Tokenize and POS-tag;
4. Remove stopwords: determiners, conjunctions, prepositions and punctuation;
5. Lemmatize words;
6. Generate n-grams;
7. Expand the n-gram set, by including references n-grams;
8. Associate article number and references;
9. Store the model.

Except for step 4, each step is responsible for adding new information to the model. The information is obtained either from the text, e.g., section title, references, or from morphological analysis, e.g., POS-tags, lemmas. If an article have references, its n-gram set is expanded with the references' n-grams. This is done so that all the necessary information for interpretation of any single article is self-contained. Besides the n-grams, links between the articles are also stored. To include the training data information, the same process is repeated for the questions, and n-gram sets from the trained questions are used to expand the associated articles' n-gram models. Since COLIEE disallowed explicit expert knowledge input, an optional information source was added after the competition, as a way of including expert knowledge in the model when available, and possibly improve system performance. This source consists in a simple term dictionary, where legal terms are associated with other correlated ones. If a given question contains n-grams referred in the dictionary, its n-gram model is expanded with the associated entries. The dictionary was written manually and contains 26 entries that were considered important after analyzing the training data, e.g., "for a third party" $\rightarrow$ "to others", and extrapolating answers to user defined queries. Tokenization and lemmatization were done using *NLTK*[5] (v. 3.0.2) with the Punkt tokenizer and WordNetLemmatizer modules, respectively.

[5] www.nltk.org

Those modules were used with their unchanged default models and settings, trained with the Punk corpus and WordNet, respectively. POS-tagging was done using *Stanford Tagger*[6] (v. 3.5.2), using the unchanged *english-left3words-distsim* model, which is trained on the part-of-speech tagged WSJ section of the Penn Treebank corpus.

To determine the relative relevance of an article with regard to the content of a question, a ranking approach was adopted. First, the n-gram set of the question is obtained by applying steps 1-6, using the question content instead of article. Then, for each article in the Civil Code, a relevance score is calculated using the following formula:

$$score = \frac{\sum_{\forall t} idf(t)}{I_q \times |q_ng_set| + I_{art} \times |art_ng_set|}, \quad t \in (q_ng_set \cap art_ng_set) \quad (1)$$

where q_ng_set is the set of n-grams for the question, art_ng_set is the set of n-grams for the article in the stored model, I_q is the relative significance of the question n-gram set size and I_{art} is the relative significance of the article n-gram set size. $idf(t)$ is the Inverse Document Frequency for the term t over the articles collection

$$idf(t) = log\frac{N}{df_t} \quad (2)$$

where N is the total number of articles and df_t is the number of articles in which t appears.

The formula (1) is a variation of the traditional *TF-IDF* scoring method, disregarding term frequency and giving different weights for the two types of document being evaluated: articles and questions, according to their size. I_q and I_{art} are parameters to be adjusted according to the corpus characteristics. This formula was developed during the first stages of analysis on the Civil Code corpus, when experiments with a TF-IDF based classifier showed poor results for this task and further observation showed that TF did not contribute for article relevance in many cases. As TF is absent from the formula, document size becomes a more relevant feature and must be considered in the scoring. In the studied corpus, law articles are usually much larger than questions in number of words, hence the different weights to adjust normalization of the score regarding the respective sets.

From this point, the articles are sorted by descending score and the 10 best are selected for filtering. The filtering step consists in fetching the best scoring article and verifying if its score exceeds a parameter threshold $confidence_thresh$. If it does, all the articles in the list that are referred by the first and exceed a parameter threshold $reference_thresh$ are also fetched. The fetched articles compose the final list of relevant articles to the input question. Parameter adjustment is described in Section 5.

[6] nlp.stanford.edu/software/tagger.shtml

4.2 Textual Entailment

A textual entailment (TE) relation in law domain comprises two levels of information. The first level describes whether or not (YES/NO, respectively) the textual evidence addresses the hypothesis. The second level describes whether the evidence supports or opposes (YES/NO, respectively) the hypothesis. However, due to the time constraint of the competition, only the first level is explored. Therefore, semantic relations such as negation and antonym were not considered in the TE evaluation step.

To detect a TE relation on a pair (Q, a), a similarity-based approach [4] can be used, in which a can answer Q if the similarity is greater than a certain threshold. However, high level inference (see Section 3) and the identification of the threshold make these methods more challenging to apply. We, therefore, propose to apply classification for detecting the TE relation with two advantages: (1) use of a rich feature set to represent data characteristics and (2) avoiding to identify the threshold.

This work shares most of the goals presented in Nguyen et. al. [7], so all the features in that work were used. However, the corpus size in this case makes it difficult to effectively train Machine Learning algorithms. For this reason, "stronger" features were sought as a way of compensating such problem. An additional Word2Vec feature was added to capture the semantic similarity of a pair (Q, a), as observation of statistical data in Table 1 shows that the lexical overlapping may not be a strong enough feature for the classification on a (Q, a) pair (e.g., cannot capture the similarity of *person* and *manager*). By adding the Word2Vec feature, the model aims to cover the semantic aspect instead of only lexical similarity. Word2Vec was trained by JPN Law corpus: a collection of all Civil law articles of Japan's constitution[7]. It contains 642 cleaned and tokenized articles, with about 13.5 million words in total.

For the classification, the Weka toolset [8] implementation of AdaBoost [18] was used, with *classifier = DecisionStump*.

Table 1: Statistical data observation in phase two

	# pairs	**# sentences**	**# tokens**	**% uni-gram word overlapping**
Training Set	267	273	36.562	58.80

The features are shown in Table 2, in which distance features measure distance between a question Q and relevant article a and statistical features capture word overlapping of this pair. After extracting features, a pipeline model was proposed and is shown in Figure 4.

In Figure 4, the first step is to preprocess the data from the input files, in which sentences and words are segmented and stopwords[9] are removed. Next, the training data is represented in a vector space model by features in Table 2. The retrieved data from relevance analysis is also denoted in the same mechanism. Finally, a classifier was trained on the training data and applied on retrieved

[7] www.japaneselawtranslation.go.jp
[8] weka.wikispaces.com
[9] https://sites.google.com/site/kevinbouge/stopwords-lists

Table 2: The feature groups; Avg is average; Q is a question, S is a sentence

	Feature	Description
Distance	Manhattan	Manhattan distance from two text fragments
	Euclidean	Euclidean distance from two text fragments
	Cosine similarity	Cosine similarity distance
	Matching coefficient	Matching coefficient of two text fragments
	Dice coefficient	Dice coefficient of two text fragments
	Jaccard	Jaccard distance of two text fragments
	Jaro	Jaro distance of two text fragments
	Damerau-Levenshtein	Damerau Levenshtein distance of two text fragments
	Levenshtein	Levenshtein distance of two text fragments
Statistical	Lcs	The longest common sub string of two text fragments
	Average of TF-IDF	Term frequency-inverse document frequency
	Avg-TF of Q and S	Avg-TF of words in a Q appearing in a S
	Avg-TF of S and Q	Avg-TF of words in a S appearing in a S
	Word overlapping	# word overlapping in a Q appearing in a article
	Average of Word2Vec	Average of word2vec similarity

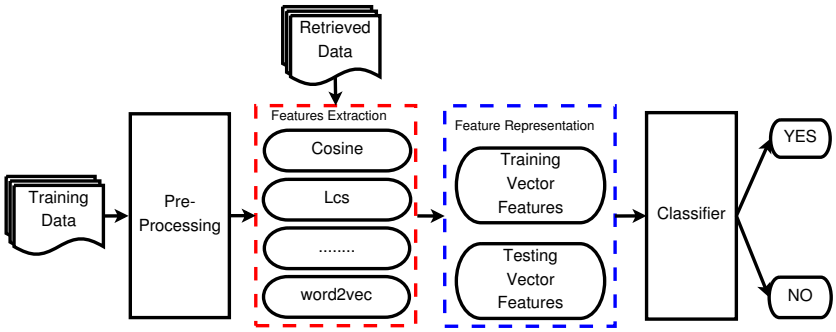

Fig. 4: The process of legal textual entailment recognition

data to judge the entailment relation. Note that features in Section 4.1 can be also used for this task.

5 Experiments and Results

5.1 Experimental Setup

The dataset was obtained from the published data for the COLIEE shared task [10], consisting in a text file with the Japanese Civil Code and a set of XML files with training and testing data for phases one to three. The training set for the three tasks contains 267 pairs (question, relevant articles). Experiments where divided in phases one and two only, dealing with Information Retrieval and Textual Entailment methods respectively. Each experiment comprised: i) data analysis, ii) model and parameter adjustments and iii) test runs.

[10] webdocs.cs.ualberta.ca/~miyoung2/COLIEE2015/

5.2 Parameter adjustment

For R_2NC, parameters I_q, I_{art}, $confidence_thresh$ (shortened to ct here), $reference_thresh$ (rt) and also k, the maximum n-gram size, were adjusted empirically on the training data using the following simple procedure:

- Starting with $I_q = 0.8$, $ct = 0.5$, $rt = 0.5$ and $k = 1$,
  1. Increase or decrease a single parameter by step $r = 0.1$ until the F-measure cannot be increased for a leave-one-out test.
  2. Repeat (1) starting from the last obtained value, with $r = 0.01$.
  3. Repeat (1) and (2) for all parameters.

For k, step was fixed on $r = 1$. I_q and I_{art} respect the constraint $I_q + I_{art} = 1$. The parameters are changed in a specific order: 1. $confidence_thesh$, 2. k, 3. $reference_thresh$, 4. I_q. I_q and I_{art} respect the constraint $I_q + I_{art} = 1$ Performance metrics were recorded for the parameter adjustment during the experiments. Fig. 5 shows the performance progression on post-competition experiments for the parameters I_q, I_{art} , with the other ones locked into their best respective values. Performance for $k \neq 3$ is negatively affected in both directions (-,+), and no further investigation was conducted for a larger range of values.

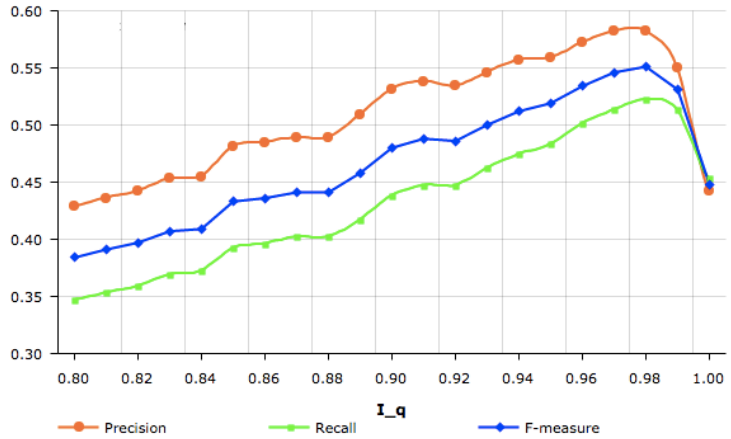

Fig. 5: Performance metrics for phase 1 related to the variation of I_q. $I_{art} = 1 - I_q$.

Final parameter values used in the competition are $k = 3$, $I_q = 0.965$, $I_{art} = 0.035$, $confidence_thresh = 0.32$ and $reference_thresh = 0.2$.

For the RTE classifier, default parameters from the Weka toolset[11] were used for all the experiments and were not changed. The parameter values are: $iterations = 10$, $seed = 1$, no *re-sampling* and $weightthreshold = 100$.

5.3 Baselines

As for the second edition of COLIEE, there is still no definite baseline for the competition dataset. However, common baseline practices and related works

[11] weka.wikispaces.com

could be used for evaluating performance on each task. For phase one, a relationship can be drawn between R_2NC and TPP [3]. For the TE task, the following baselines were used for comparison:

- SVMs: uses Support Vector Machines (SVMs)[12] [19] with Weka. The parameters are $C = 1$, $\gamma = 0$, *kernel Type = radial basis function (RBF).*
- AdaBoost-SVMs: uses SVMs as weak learners instead of DecisionStump.

5.4 Evaluation Method

Given the limited training data available, leave-one-out validation was used to evaluate the performance of the model in both tasks on the training dataset with three measures: precision (P), recall (R) and F-measure (F) as in Eq. (3), (4) and (5). In phase two, accuracy (A) measurement is also used as in Eq. (6).

$$P = \frac{Cr}{Rt} \quad (3) \qquad R = \frac{Cr}{Rl} \quad (4) \qquad F = \frac{2(P * R)}{P + R} \quad (5) \qquad A = \frac{Cq}{Q} \quad (6)$$

where Cr counts the correctly retrieved articles for all queries, Rt counts the retrieved articles for all queries, Rl counts the relevant articles for all queries, Cq counts the queries correctly confirmed as true or false and Q counts all the queries.

5.5 Pre-competition Results

Pre-competition experiment results on the shared data are presented in Tables 3 and 4.

Table 3: Experiment results for phase one (IR) with R_2NC. In the top 3/10 settings, articles ranked up to 3^{rd} or 10^{th} place are marked as relevant.

	Precision	**Recall**	**F-measure**
R2NC	0.568	0.516	**0.54**
R2NC (top 3)	0.27	0.64	0.38
R2NC (top 10)	0.10	0.77	0.17
TPP (top 3)	N/A	0.52	N/A
TPP (top 10)	N/A	0.91	N/A

Table 4: AdaBoost-DecSt (*DecisionStump*) vs. SVMs and AdaBoost-SVMs.

	Precision	**Recall**	**F-measure**	**Accuracy (%)**
AdaBoost-DecSt	0.621	0.614	0.617	**61.42**
SVMs	0.537	0.543	0.539	54.30
AdaBoost-SVMs	0.485	0.491	0.487	49.06

The results indicate that R_2NC is expected to be competitive with state-of-the-art approaches to relevance analysis in legal documents, such as *TPP* [3]. However, the proposed method is much simpler when compared to *TPP* and operates with considerably less training data: 266 documents for R_2NC against 1518 documents for *TPP*. R_2NC design also makes it difficult for the model to be overtrained beyond the parameter adjustment, since no training data is counted

[12] https://www.csie.ntu.edu.tw/∼cjlin/libsvm/

more than one time and the method is single-shot, as opposed to convergence-based. Experiments were repeated with traditional TF-IDF scoring instead of R_2NC formula, yielding 0.51 F-measure.

Results of RTE in Tab. 4 indicate that AdaBoost with a set of appropriate features outperforms the baselines by 7.74% (*SVMs*) and 12.94% (*Adaboost-SVMs*) on F-measure. Moreover, the precision and accuracy of this method also achieve considerable improvements when compared to the baselines. This suggests that the features are expected to be efficient for addressing TE in the legal domain. This conclusion is supported by the accuracy measurements.

Another interesting point is that Word2Vec similarity contributes to improve the performance of RTE. As stated in Section 3, legal documents usually require concept linking to understand and answer a question; therefore, semantic similarity from Word2Vec helps to improve the performance. The results also show the efficiency of the lexical features.

The performance of RTE in the law domain, however, is not comparable with the same task in common data i.e., news articles [6,7] due to the characteristics of law dataset, as shown in Section 3. The performance was not improved very much even when many features in both phase one and two were combined. This suggests that more sophisticated approaches e.g., semantic inference or semantic rules should be considered in feature construction. Finally, negation and antonym analysis should be considered to improve the quality of the entailment recognition, effectively exploring the second level of entailment information as described in Section 4.2.

5.6 Feature Evaluation

Further evaluation of feature impact on TE model was conducted by leave-one-out test. The most effective features are shown in Table 5.

Table 5: Top 4 influential features, italic is for statistical features. Values are the difference in F-measure between the model with all features and without the single specified feature.

Features	**Influential value**	**Features**	**Influential value**
Euclidean	0.005	*Lcs*	0.0001
Damerau-Levenshtein	0.154	Average of Word2Vec	0.024

Table 5 shows an indication of contribution from features to the model. Results show that all effective features contribute to the method. Note that both *Damerau-Levenshtein* and *Euclidean* are distance features whereas *the longest common substring* (lcs) is a statistical feature. The results support that in legal texts, there is not much word overlapping between a question and relevant articles. An interesting aspect is that Word2Vec similarity has a big positive impact to the model. This supports the conclusion on similarity stated in Section 5.5.

5.7 Competition Results

The method presented in this paper achieved significant results in COLIEE, being ranked 2^{nd} in phase one (IR) and 3^{rd} in phase three (combined IR + TE). It was not well ranked in phase two (TE). The relevant competition results are presented in Table 6 as they were announced in JURISIN 2015.

Table 6: Competition results for phases 1 (IR) and 3 (IR + TE) respectively. First three ranked.

Rank	ID	Prec.	Recall	F-m
1	UA1	0.633	0.490	0.552
2	**JAIST1**	**0.566**	**0.460**	**0.508**
3	ALV2015	0.342	0.529	0.415

Rank	ID	Accuracy.
1	UA1	0.658
2	Kanolab3	0.620
3	**JAIST1**	**0.582**

5.8 Post-competition Analysis and Improvements

Post competition analysis pointed us to possible sources of classification problems in phase 2 (TE) and also gave directions of improvement in both tasks.

For R_2NC, the lack of an implicit semantic mapping was an important factor when compared to the top ranked approach. To compensate for that, a term dictionary was included as a new information source for expanding the question n-gram models as described in Section 4.1. By using linguistic observations, it was possible to create basic entries in the dictionary (non-expert knowledge), improving phase 1 F-measure on the shared data (Table 3) from 0.54 to 0.55.

In the case of phase 2, over-fitting on training data was deemed the main factor that reduced classification performance. Our system achieved over 61% accuracy (Table 4) when running on the shared data, but only 37.88% reported from the competition results. Phase three results show that accuracy improved when restricting information for the classifier and this is consistent with the over-fitting assumption. Another important point is that a question q and all sentences in an article a were used in building the vector space model. As a result, imbalance of length between the question and the article may have affected feature calculation. This can be addressed by developing a better text segmentation method. Finally, the over-fitting assumption can also be dealt by using other classification approaches e.g., Deep Neural Networks, together with over-fitting avoidance techniques e.g., pruning, dropout.

5.9 Error Analysis and Discussion

An investigation was done on the ranked list obtained with R2NC in phase one (see Section 4.1). It revealed that relevant articles ranked 3rd and below had keywords that did not appear in the corresponding question in the corpus. This reinforces the view that the questions are highly directed, albeit in a conceptual level. Relevant articles that ranked lower than 15th (approx. 20%) were found to require a relatively high level of abstraction to obtain an interpretation that could link to the corresponding question. Table 7 shows an example of complex relevance relationship.

Table 8 shows a case in which our system gives correct outputs (ID H18-2-4). In this example, there are several common words from which this approach can correctly judge the TE relation, e.g., *reimbursement*. In addition, several words can be inferred from the questions by using Word2Vec similarity e.g., *person* $\sim$ *manager*, *fees* $\sim$ *costs or expenses*. This supports our observation that TE can be addressed by using lexical features and word similarity. For example, in (ID H18-2-4), our system can still predict the TE relation correctly, even with little

Table 7: Example of pair (question, article) with low ranking but high relevance.

ID	Article	Question	Ranked in
H18-2-1	Article 697(1)A person who commences the management of a business for another person without being obligated to do so (hereinafter in this Chapter referred to as "Manager") must manage that business (hereinafter referred to as "Management of Business") in accordance with the nature of the business, using the method that best conforms to the interests of that another person (the principal).(2)The Manager must engage in Management of Business in accordance with the intentions of the principal if the Manager knows, or is able to conjecture that intention.	In cases where a person plans to prevent crime in their own house by fixing the fence of a neighboring house, that person is found as having intent towards the other person.	424th

Table 8: Examples of entailment judgment; P is predicted and A is annotated

ID	Article	Question	P	A
H18-2-4	(Managers' Claims for Reimbursement of Costs)Article 702(1)If a Manager has incurred useful expenses for a principal, the Manager may claim reimbursement of those costs from the principal.(2)The provisions of Paragraph 2 of Article 650 shall apply mutatis mutandis to cases where a Manager has incurred useful obligations on behalf of the principal.(3)If a Manager has engaged in the Management of Business against the intention of the principal, the provisions of the preceding two paragraphs shall apply mutatis mutandis, solely to the extent the principal is actually enriched.	In cases where a person repairs the fence of a neighboring house after it collapsed due to a typhoon, but the neighbor had intended to replace the fence with a concrete-block wall in the near future, if a separate typhoon causes the repaired sections to collapse the following week, reimbursement of repair fees can no longer be demanded.	YES	YES
H18-26-1	(Renunciation of Shares and Death of Co-owners)Article 255 If one of co-owners renounces his/her share or dies without an heir, his/her share shall vest in other co-owners.	In cases where person A and person B co-own building X at a ratio of 1:1, if person A dies and had no heirs or persons with special connection, ownership of building X belongs to person B.	NO	YES

lexical overlap. This indicates the efficiency of this approach, and especially of the word similarity feature.

On the other hand, the pair H18-26-1 exemplifies a case in which the system predicted NO while TE relation was annotated YES even when the question and answer share more common words. This shows the limitation of this feature set in cases where the question and answer are short. In this case, after removing stop words, a few remaining words may not be enough to capture the TE relation. Moreover, the lack of important words e.g., *building*, *connection* or *belong* reveals a big challenge for our system to decide the TE relation. This suggests that a keyword enriching mechanism such as term expansion used in phase one could improve the results.

In order to facilitate the understanding of different error cases and give other people the opportunity to try the system developed for the competition, an online demo system[13] has been made available. In this demo it is possible to input user defined questions or just verify the answers to questions in the COLIEE shared data.

[13] http://150.65.242.101:3001/

6 Conclusion

This paper explores the challenging issue of building a QA system in the legal domain. We propose a model including three stages: legal information retrieval, legal textual entailment and legal text answering. In the first stage, a mixed size n-gram model built from morphological analysis is used to rank and select relevant articles corresponding to a legal question; next, pairs of questions and retrieved articles are judged by a machine learning algorithm trained on lexical features and Distributional Semantic similarity, to decide whether the questions can be answered positively or negatively by the retrieved articles; and finally, correct answers would be provided for users in the final stage. The contributions of this work in IR and TE task are: 1) a simple, yet effective language model for law corpora coupled with a Relevance Analysis method (R_2NC) capable of exploiting such model; 2) The use of TF-IDF and Word2Vec similarity features for applying Machine Learning algorithms to RTE. With a recall of 0.64 for the top 3 ranked articles, R_2NC appears as competitive when compared to state-of-the-art similar work, in spite of being more simple and applicable with less training data. By combining lexical features and Word2Vec similarity, this approach for LQA also outperformed the baselines by 8.4% (SVM) and 11.3% (*Adaboost-SVMs*) on F-measure. Results in the COLIEE competition for the IR task (0.508 F-measure, 2^{nd} place) and the combined IR+TE task (0.582 accuracy, 3^{rd} place) indicate a substantial adequacy to the tasks addressed. The competition also provided important shortcomings of the proposed approach, namely the lack of implicit semantic representation and classifier over-fitting. Those shall be addressed in future work.

Still on future directions, information on a higher abstraction level, e.g., syntactic mappings, could be used to improve the language model for the IR task. In the TE task, since a sentence in a legal article is usually long, a sophisticated method of sentence partition e.g., requisite and effectuation should be considered. In feature extraction, features in IR should be combined with lexical features in TE and investigated to improve the quality of the judgment. Moreover, capturing contradictions in the TE relation by current statistical features is a big challenge. To solve this issue, semantic rules over negation and antonym detection should be defined and incorporated into the feature extraction. Finally, we would like to investigate and apply sentence similarity calculation by Sent2Vec to improve the performance of the TE.

Acknowledgements

This work is supported partly by the grant of NII Research Cooperation and JAIST's Research grant.

References

1. Robert C. Berring: "The heart of legal information: The crumbling infrastructure of legal research". *Legal information and the development of American law*. St. Paul, MN: Thomson/West, 2008.

2. Rinke Hoekstra, Joost Breuker, Marcello Di Bello and Alexander Boer: "The LKIF Core ontology of basic legal concepts". Proc. of the Workshop on Legal Ontologies and Artificial Intelligence Techniques (LOAIT 2007), 2007.
3. Yi-Hung Liu, Yen-Liang Chen and Wu-Liang Ho: "Predicting associated statutes for legal problems". *Information Processing & Management* 51.1: 194-211, 2015.
4. Quang-Thuy Ha, Thi-Oanh Ha, Thi-Dung Nguyen, and Thuy-Linh Nguyen Thi: "Refining the judgment threshold to improve recognizing textual entailment using similarity." Computational Collective Intelligence. Technologies and Applications. Springer Berlin Heidelberg: 335-344, 2012.
5. Diana Inkpen, Darren Kipp and Vivi Nastase: "Machine Learning Experiments for Textual Entailment. *Proceedings of the Second Challenge Workshop Recognising Textual Entailment*: 17-20, 2006.
6. Julio Javier Castillo: "An approach to Recognizing Textual Entailment and TE Search Task using SVM. *Procesamiento del Lenguaje Natural*, 44: 139-145, 2010.
7. Minh-Tien Nguyen, Quang-Thuy Ha, Thi-Dung Nguyen, Tri-Thanh Nguyen and Le-Minh Nguyen: "Recognizing Textual Entailment in Vietnamese Text: An Experiment Study." KSE: 108-113 2015.
8. Quang Nhat Minh Pham, Le Minh Nguyen and Akira Shimazu: "Using Machine Translation for Recognizing Textual Entailment in Vietnamese Language." RIVF: 1-6, 2012.
9. Danilo Giampiccolo, Bernardo Magnini, Ido Dagan and Bill Dolan: "The third PASCAL recognising textual entailment challenge". *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*. Association for Computational Linguistics: 1-9, 2007.
10. Oanh Thi Tran, Bach Xuan Ngo, Minh Le Nguyen and Akira Shimazu: "Answering Legal Questions by Mining Reference Information". New Frontiers in Artificial Intelligence. Springer International Publishing: 214-229, 2014.
11. Mi-Young Kim, Ying Xu, Randy Goebel and Ken Satoh: "Answering Yes/No Questions in Legal Bar Exams". New Frontiers in Artificial Intelligence. Springer International Publishing: 199-213, 2014.
12. Bach Xuan Ngo, Minh Le Nguyen and Akira Shimazu: "RRE Task: The Task of Recognition of Requisite Part and Effectuation Part in Law Sentences." J. IJCPOL 23(2): 109-130, 2010.
13. Oanh Thi Tran, Bach Xuan Ngo, Minh Le Nguyen and Akira Shimazu: "Reference Resolution in Legal Texts." In Proc. of ICAIL: 101-110, 2013.
14. Ido Dagan, Bill Dolan, Bernardo Magnini and Dan Roth: "Recognizing textual entailment: Rational, evaluation and approaches - Erratum". *Natural Language Engineering* 16(1): 105-105, 2010.
15. Rui Wang: *"Intrinsic and Extrinsic Approaches to Recognizing Textual Entailment.* Saarland University, ISBN 978-3-933218-32-2: 1-219, 2011.
16. Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn and Tony Robinson: "One billion word benchmark for measuring progress in statistical language modeling", arXiv preprint arXiv:1312.3005, 2013.
17. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado and Jeffrey Dean: "Distributed Representations of Words and Phrases and their Compositionality". In Proceedings of NIPS, 2013.
18. Yoav Freund and Robert E. Schapire: "A decision-theoretic generalization of online learning and an application to boosting". Journal of computer and system sciences 55.1: 119-139, 1997.
19. Corinna Cortes and Vladimir Vapnik: "Support-Vector Networks". *Machine Learning* 20(3): 273-297, 1995.